

REMARKS

This Amendment is responsive to the Final Office Action mailed June 5, 2008, hereinafter referred to as “the Office Action.”

Claims 1-7 and 32-43 remain pending and new claims 49-58 are added by this Amendment. Applicants note with appreciation the withdrawal of all previous rejections. However, the Office Action sets forth new grounds of rejection, which are addressed below.

Amendments

Claims 44-48 are canceled by this Amendment without prejudice. Significant amendments are made to the claims and Applicants provide, for the convenience of the Examiner, a clean copy of the claims as amended in the attached Appendix. Independent claims 1 and 37 are significantly amended to recapture the concept of selecting a path based at least in part on contents of multipath routing information and VM-specific information. Applicants have reconsidered the Saito reference as explained in more detail below, and believe claims 1 and 37 are patentable over Saito and secondary-cited references.

Claims 1 and 37 are further amended to clarify that the request to transfer data is actually a request to transfer data between the VM and a virtual device backed up by a data storage unit. As explained in the written description (e.g., at paragraph 72 or 123), virtualization software emulates devices for the VM, and virtualizes these devices using hardware resources, essentially backing up these resources.

The claims are also amended to include the concept of an algorithm that determines whether to perform an action (in the case of claim 37, suspending the first VM, and in the case of claim 44, migrating the first VM). The concept of an algorithm is explained in the specification as filed, e.g., at paragraphs 56-57 and elsewhere.

Claims 37 is further amended to clarify that the computer program, when executed on the computer system, causes the computer system to implement a method.

Clarifying amendments are made to claim 32 and claims depending therefrom.

New claims are provided to capture subject matter previously presented in claims that are now canceled.

No new matter is introduced by this Amendment.

Specification

1. Objection of the Specification under 37 C.F.R. § 1.75(d)(1)

The specification is under objection for “failing to provide a proper antecedent basis for the claimed subject matter.” In particular, the Office Action states that “claims 37-48 are rejected for reciting a ‘computer program’ and a ‘computer-readable medium’ but Applicant’s specification fails to provide any description for these terms” (Office Action, page 2¹).

Applicants respectfully traverse because the terms are well known terms of art, and are clear and definite when read in light of the specification as a whole. Furthermore, Applicants respectfully note that no rejection under 35 U.S.C. §112 second paragraph were made against the claims containing the allegedly objectionable phrases, which suggests that the claims and these terms are, in fact, definite.

The Office Action relies on 37 C.F.R. § 1.75(d)(1) and MPEP § 608.01(o) in making this objection. 37 C.F.R. § 1.75(d)(1) states, “The claim or claims must conform to the invention as set forth in the remainder of the specification and the terms and phrases used in the claims must find clear support or antecedent basis in the description so that the meaning of the terms in the claims may be ascertainable by reference to the description” (emphasis added). MPEP 608.01(o) likewise states that “The meaning of every term used in any of the claims should be apparent from the descriptive portion of the specification with clear disclosure as to its import; and in mechanical cases, it should be identified in the descriptive portion of the specification by reference to the drawing, designating the part or parts therein to which the term applies” (emphasis added). It is noted that this Rule and this MPEP section are directed to situations where a claim term is unclear and whose meaning cannot be ascertained by review of the specification. They do *not* contain an explicit requirement that each term or phrase used in the

¹ The pages of the Office Action were not numbered. References to the pages of the Office Action are therefore made consistent with normal PTO practice of numbering the first page after the “Office Action Summary” form (PTOL-325) as page 2, and then consecutively thereafter.

claims be matched with an antecedent in the specification, only that such terms and phrases “find clear support . . . so that the meaning of the terms in the claims are ascertainable by reference to the description.”

However, the Office Action does not anywhere suggest that the meanings of the terms are unclear. For instance, The Office Action does not reject claims 37-48 for being indefinite under 35 U.S.C. §112, nor is the specification objected to under 35 U.S.C. §112 for lacking support for these terms under the Written Description requirement, which suggests acknowledgement that Applicants indeed had possession of the concepts of “computer program” and “computer-readable medium” at the time the application was filed. Since the meanings of the terms are clear, Applicants respectfully submit that objection to the specification under 37 C.F.R. § 1.75(d)(1) is inappropriate.

As stated in MPEP 2173.05(e), “The mere fact that a term or phrase used in the claim has no antecedent basis in the specification disclosure does not mean, necessarily, that the term or phrase is indefinite. There is no requirement that the words in the claim must match those used in the specification disclosure. Applicants are given a great deal of latitude in how they choose to define their invention so long as the terms and phrases used define the invention with a reasonable degree of clarity and precision” (emphasis added).

Applicants accordingly respectfully submit that the terms, “computer program” and “computer-readable medium,” being well known terms in the field of computer software, to which the present application is directed, meet the test of “reasonable clarity and precision” and are sufficiently supported by the specification, although these precise words are not used. Reconsideration and retraction of this objection is respectfully requested.

2. Requirement of correction of claims 37-48.

The Office Action does not cite any authority to require Applicants to “correct” claims 37-48. Applicants respectfully traverse and submit that that which is not broken need not be fixed. The claims are not rejected under 35 U.S.C. §112, second paragraph; hence, they are not indefinite. The Office Action contains no allegation that the specification fails to meet the Written Description or Enablement requirement under 35 U.S.C. §112 first paragraph for these terms; hence, the specification sufficiently enables and shows possession of these concepts by the

inventors at the time the application was filed. Although the precise terms “computer program” and “computer-readable medium” are not identically recited in the written description, these are well known terms of art, the meanings of which are readily ascertainable upon reading the specification. As noted above, there is no requirement in U.S. patent law or regulations that would compel that a specification contain dictionary definitions of every term or phrase used in every claim, as acknowledged in MPEP 2173.05(e), quoted above. It is ironic that this entire application is directed to the fields of computer programs and data storage, and yet the terms, “computer program” and “computer-readable medium” are assailed because, “Applicant’s specification fails to provide any description for these terms” (Office Action, page 2). For the above reasons, reconsideration and retraction of this requirement is respectfully requested.

Claim Objections

Claims 38-43 and 45-48 are under objection for being improperly dependent under 37 C.F.R. § 1.75(c). Applicants, respectfully, do not understand this objection. It appears that the Office Action is suggesting that claims depending from independent claims 37 and 44 are improper because they don’t directly refer in the body of the claims to the computer program set forth in those independent claims, and that somehow this means that the claims do not incorporate all the limitations of depended-upon claims. Applicants respectfully submit that this is an incorrect view of the claims.

With regard to incorporation by dependent claims of the limitations set forth in depended-upon claims, the Rules are very clear: “Claims in dependent form shall be construed to include all the limitations of the claim incorporated by reference into the dependent claim” (37 C.F.R. § 1.75(c)) (emphasis added). Therefore, since each of claims 38-43 depend from claim 37, they necessarily incorporate all the limitations of the claim 37 (as well as any intervening claim in the chain of dependency). Furthermore, it is perfectly appropriate for a dependent claim to further define a single element of a depended-upon claim, without having to recite what feature that element is a part of, and so on. Applicants respectfully request reconsideration and withdrawal of this objection.

Claim Rejections - 35 U.S.C. §112 first paragraph

Claims 3, 34, and 39 stand rejected under 35 U.S.C. §112, first paragraph for failing to comply with the Written Description requirement. Specifically, the Office Action (page 3-4) suggests that the specification lacks support for a feature of not routing data to the data storage unit. The Office Action further notes that, “This rejection may be overcome if Applicant can cite specific sections in the specification that support the subject matter being claimed.”

Applicants respectfully traverse the rejection. Applicants appreciate this opportunity to identify how this concept is supported by the specification. Referring to Figure 7, at step 810, it is determined, based on VM-specific information and path information, whether to Route or Queue a specific data transfer request. This decision is described in paragraphs 108-114. Then in paragraph 122, the specification describes the decision as to whether *or not* to route queued requests. As stated in the specification, “The decision whether or not to route queued requests at step 826 may be substantially the same as the decision at the step 810.”

For the above-stated reasons, Applicants respectfully submit that the specification adequately supports the concept of not routing data to the storage unit, and accordingly requests that this rejection be withdrawn.

Claim Rejections - 35 U.S.C. § 103(a)

All claims 1-7 and 32-48 stand rejected under 35 U.S.C. § 103(a) for being unpatentable over U.S. Patent 5,257,386 issued to Saito (hereinafter, “Saito”) in view of U.S. Patent 7,213,246 issued to van Rietschote (hereinafter “Rietschote”). Applicants note that claims 44-48 are canceled by this Amendment without prejudiced. However, with respect to the claims not canceled that are rejected, Applicants respectfully traverse because the prior art fails to teach or suggest each and every feature set forth in the claim, and because the prior art lacks motivation to combine and/or modify the references as proposed by the Office Action.

Claims 1, 37

Claims 1 and 37 set forth, *inter alia*:

“selecting ne path of the plurality of paths according to an algorithm, which takes as inputs at least contents of the multipath routing information and contents of the VM-specific information;

and

“routing a physical data transfer request corresponding to the virtual data transfer request to the data storage system over the one path that was selected”

Applicants respectfully submit that at least this feature is not taught or suggested by Saito, nor any reference of record. This feature is similar to that claimed in the application as originally filed, which was addressed in the non-final Office Action mailed August 24, 2007. That Office Action stated that “Saito discloses . . . based on multipath routing information and the VM-specific information, deciding whether to route the data transfer request [column 3 <<line67>> to column 4 <<line 14>> | column 6 <<lines 36-46>>: decision on whether to route the request based on whether the path is available and the VM’s transfer priority]” Upon reconsideration of Saito, Applicants respectfully disagree with the Examiner’s characterization of Saito.

The first identified portion of Saito, col. 3, line 67 to col. 4 line 14, is directed a component of Saito called a “transfer length limiter” identified in Figure 1 at 15. This component is not coupled to transfer paths 20. Instead, limiter 15 is responsible for limiting transfer throughput based on the transfer priority of the source VM. If a requested transfer data length is longer than an imposed limit, the transfer request can be split into two. The first portion can be sent immediately and the second portion sent back to queue 13, thereby preventing a single VM from monopolizing storage bandwidth. Col. 4, lines 15-17 explain that the transfer instruction substituting part issues the transfer request to transfer data between the main storage 19 and the external storage. There is no suggestion in the indicated portion of selecting a path based on VM-specific information.

The second identified portion of Saito, col. 6, lines 36-46 addresses path selection, and specifies that the path selector “selects a free path 20 between the main storage 19 and the external storage 18 for the request.” Thus, Saito supports selecting a path based on multipath routing information, but does not suggest selecting a path based on VM-specific information.

Therefore, while Applicants’ system, for example, supports the possibility of reserving one or more paths to a specific VM, there does not appear to be any suggestion that Saito can do the same. The secondary reference Rietschote and other cited references fail to overcome the deficiencies of Saito. Since independent claims 1 and 37 include features that are not taught or

suggested by the prior art of record, Applicants respectfully submit that they are patentable and should be allowed. Claims depending from claims 1 and 37 further distinguish the invention from the prior art and are therefore are likewise allowable for at least the same reasons as claims 1 and 37.

Claims 2, 3, 5-7, 32-36, 38, 41-43, 53 and 55

Claims 2, 3, 5-7, 32-36, 38, 41-43, 53 and 55 are directed, in general, to an action that can be taken upon the failure of one of the data paths communicating data between the computer system and the data storage. Certain claims include suspending a VM in response to a data path failure and other claims include migrating a VM in response to a data path failure. The Office Action cites Rietschote for teaching fail-over in virtualized computer systems. Specifically, the Office Action states:

“Rietschote discloses determining when an application has failed in a cluster of computer systems. . . . Clearly, since these applications are well know to be involved in the transfer of data of a connection path, when these types of applications fail, transfer of data from these applications over the network are prevented. . . . Therefore, Rietschote discloses the missing limitations of determining that a failure has occurred that prevents the transfer of data over a first path of the plurality of paths and subsequently suspending the VM in response to determining that a failure has occurred”

(Office Action, page 5). Applicants do not agree that one can take performing management actions on a VM (e.g., suspending a VM) based on the failure of an application and extrapolating from that to performing management actions based on the failure of a data path in a multipath data storage system.

Initially, Applicants agree that Saito does not contemplate failures in the storage data path. Applicants further agree that Rietschote discusses suspending a VM in response to a failure in an application running on the VM. Applicants also concede that detection and remediation of network failovers was known. (See, specifically, paragraphs 45-47 of the Application as filed.) However, it was not known to perform a VM management action in response to a failure in a storage path, and Applicants respectfully submit that it would not have been obvious at the time the invention was made to do so.

Under *KSR*², a number of rationales are available to the Patent Office to issue an obviousness rejection. However, regardless as to the rationale, it must be clearly articulated why the claimed invention would have been obvious. Applicants respectfully submit that such a clear articulation has not been provided. For instance, it has not been explained how the two references, Saito and Rietschote, could be combined to arrive at the claimed invention or why such a combination would have been attempted.

Specifically, Applicants note that, in order to take a VM management action (e.g., suspending a VM) in response to a storage data path failure, information about the failure must be provided to the hypervisor kernel. None of the references suggest an architecture that enables such an action to be taken. As explained in the present Application with reference to Figures 3 and 4, in traditional implementations of a multipath storage architecture, and in reference to Figure 4 of the present Application, storage path manager 38C is responsible for managing multipath routing information, and this information is not readily available to kernel 600B. As stated in the application as filed, “the SPM 38C would not have much information about the VMs 200 to 200N or the VMMs 300 to 300N, and the kernel 600B would not have much information about the routing of data by SPM 38C” (paragraph 73). Without having multipath routing information available at the kernel, it would not be possible to take VM management actions, such as whether to suspend a VM, based on the multipath routing information. As explained in the Application with reference to Figure 5, Applicants have solved this problem by integrating the storage path manager 642 with kernel 600C. Because software components are structured in a particular way in prior art implementations, e.g., as described in the present Application with reference to Figure 3, the prior art systems are not capable of performing failover actions as described in claims 2, 3, 5-7, 32-36, 38, 41-43, 53 and 55.

To Applicants’ knowledge, Applicants’ were first to discover that a failure in a storage data path could disrupt VM execution due to congestion of remaining storage data paths, which ultimately led to the invention. Nor are Applicants aware of any other motivation to perform the operations described in the above-referenced claims. Furthermore, it is one thing to suspend a faulty VM, i.e., a specific VM experiencing a failure of software running thereon, as suggested by Rietschote. It is another thing entirely to identify a low priority VM of a plurality of VMs

² *KSR International Co. v. Teleflex Inc. (KSR)*, 550 U.S. ___, 82 USPQ2d 1385 (2007)

and, based on the priority of that VM, suspend it when a storage network failure occurs that potentially impacts all VMs. Thus, arriving at Applicants' claimed invention requires more than simply substituting one known element for another, or applying a known technique to solve a known problem.

For the reasons set forth above, Applicants respectfully submit that the Office Action fails to clearly articulate a reasoning or rationale for combining the references Saito and Rietschote to arrive at subject matter of claims 1, 2, 5-7, 32-36, 38, 41-43, 53, and 55. Instead, the Office Action simply presents two prior art references peripherally related to multipath storage and VM management along with some hand-waving to conclude that Rietschote "discloses the missing limitations." Accordingly, Applicants respectfully request reconsideration of the rejection as it relates to claims 2, 5-7, 32-36, 38, 41-43, 53, and 55.

Remaining Claims

Remaining claims not specifically addressed above are believed to be patentable for at least the same reasons described above with respect to claim upon which they depend.\

Applicants respectfully submit that the present Application is in condition for allowance. Applicants therefore respectfully request reconsideration of the outstanding rejections and a Notice of Allowance. The Examiner is invited to contact the undersigned at 650-427-2390 to discuss any additional changes the Examiner may feel is necessary in light of this Amendment.

Date: October 2, 2008

Respectfully submitted,

/Leonard Heyman/

Leonard Heyman
Reg. No. 40,418
Attorney for the Applicant

APPENDIX – CLEAN LISTING OF CLAIMS

The following listing of the claims is provided for the convenience of the Examiner and shows the claims as amended, without the mark-up.

1. (Currently amended) A method for responding to a request to transfer data between a first virtual machine (VM) in a computer system and a virtual storage device backed up by a data storage unit within a multipath data storage system, the method comprising:

identifying the request issued by the first VM, the request being a virtual data transfer request, the first VM being one of a plurality of VMs;

identifying a plurality of possible paths over which the data could be routed from the computer system to the data storage system and multipath routing information related to a state of each of the possible paths;

determining VM-specific information related to the first VM;

selecting one path of the plurality of paths according to an algorithm, which takes as inputs at least contents of the multipath routing information and contents of the VM-specific information; and

routing a physical data transfer request corresponding to the virtual data transfer request to the data storage system over the one path that was selected.

2. (Currently amended) The method of claim 1, further comprising:

determining whether a failure has occurred that prevents the transfer of data over a first path of the plurality of possible paths;

failing over to one or more alternate paths when the failure has occurred, deciding whether the first VM should be suspended according to a second algorithm, wherein the second algorithm takes as inputs at least whether the failure has occurred and contents of the VM-specific information, and suspending the first VM when the second algorithm returns a decision to suspend the first VM.

3. (Currently amended) The method of claim 2, further comprising not routing the data to the data storage unit.

4. (Currently amended) The method of claim 1, in which the VM-specific information indicates a priority of the first VM relative to other VMs of the plurality of VMs.

5. (Currently amended) The method of claim 2, wherein the VM-specific information indicates a priority of the first VM relative to other VMs of the plurality of VMs, and the second algorithm results in a decision to suspend the first VM when the first VM is determined to have a lower priority than one or more other VMs of the plurality of VMs and the failure is determined to have occurred.

6. (Currently amended) The method of claim 2, wherein the suspending of the first VM includes suspending the first VM until the failure is corrected.

7. (Currently amended) The method of claim 2, wherein the suspending of the first VM includes suspending the first VM until a failback occurs.

8-31. (Cancelled)

32. (Currently amended) A method for responding to a request to transfer data between a first virtual machine (VM) in a computer system and a virtual storage device backed up by a data storage unit within a multipath data storage system, the method comprising:

identifying the request issued by the first VM, the request being a virtual data transfer request, the first VM being one of a plurality of VMs;

identifying a plurality of possible paths over which the data could be routed from the computer system to the data storage system and multipath routing information related to each of the possible paths;

determining whether a failure has occurred that prevents the transfer of data over a first path of the plurality of possible paths;

determining VM-specific information related to the first VM; and

when the failure is determined to have occurred, deciding whether the first VM should be migrated to a different physical computer system according to an algorithm, which takes as inputs at least contents of the VM specific information, and migrating the first VM to a different physical computer when the failure has occurred.

33. (Currently amended) The method of claim 32, further comprising, when the failure is determined to have occurred, failing over to one or more alternate paths.

34. (Currently amended) The method of claim 32, further comprising not routing the data the data storage unit.

35. (Previously presented) The method of claim 32, in which the VM-specific information indicates the first VM's priority relative to other virtual machines.

36. (Previously presented) The method of claim 35, wherein the first VM is determined to have a lower priority than one or more other virtual machines.

37. (Currently amended) A tangible computer readable medium embodying a computer program for handling a data transfer request in a computer system, the computer system comprising virtualization software interposed between and interfacing with a plurality of virtual machines (VMs) and system hardware, the computer program being integrated with or

coupled to the virtualization software, whereby, when executed on the computer system, the computer program causes the virtualization software to perform a method for handling a data transfer request between a first VM of the plurality of VMs and a virtual device backed up by data storage unit within a multipath data storage system, the method comprising:

identifying the data transfer request issued by a first VM;

identifying a plurality of possible paths over which the data could be routed from the computer system to the data storage system and multipath routing information related to a state of each of the possible paths;

determining VM-specific information related to the first VM;

selecting one path of the plurality of paths according to an algorithm, which takes as inputs at least contents of the multipath routing information and contents of the VM-specific information; and

routing a physical data transfer request corresponding to the virtual data transfer request to the data storage system over the one path that was selected.

38. (Currently amended) The tangible computer readable medium of claim 37, wherein the method further comprises:

determining whether a failure has occurred that prevents the transfer of data over a first path of the plurality of possible paths;

failing over to one or more alternate paths when the failure has occurred;

deciding whether the first VM should be suspended according to a second algorithm, which takes as inputs at least whether the failure has occurred and contents of the VM-specific information; and

suspending the first VM when the second algorithm returns a decision to suspend the first VM.

39. (Currently amended) The tangible computer readable medium of claim 37, wherein the method further includes not routing the data to the data storage unit.

40. (Currently amended) The tangible computer readable medium of claim 37, in which the VM-specific information indicates a priority of the first VM relative to other VMs of the plurality of VMs.

41. (Currently amended) The tangible computer readable medium of claim 38, wherein the algorithm results in a decision to suspend the first VM when the first VM is determined to have a lower priority than one or more other VMs of the plurality of VMs and the failure is determined to have occurred.

42. (Currently amended) The tangible computer readable medium of claim 38, wherein the suspending of the first VM includes suspending the first VM at least until the failure is corrected.

43. (Currently amended) The tangible computer readable medium of claim 38, wherein the suspending of the first VM includes suspending the first VM at least until a failback occurs.

44-48. (Canceled)

49. (New) The method of claim 1, further comprising:
deciding, prior to the routing and according to a second algorithm, whether to immediately route the request or to queue the request, the second algorithm taking as inputs at least contents of the multipath routing information and contents of the VM-specific information,

wherein the routing is immediately performed when the deciding results in a decision to immediately route the request.

50. (New) The method of claim 1, wherein the VM-specific information indicates an amount of disk bandwidth that is allocated to the VM.

51. (New) The method of claim 1, wherein the multipath routing information includes a pending data transfer load for each of plurality of possible paths over which data could be routed.

52. (New) The method of claim 51, wherein:

the VM-specific information includes an identifier of the first VM as the source of the request; and

the algorithm distributes requests such that substantially all requests received from the first VM are routed over the one path, and substantially all requests from at least another VM of the plurality of VMs are routed over a second path of the plurality of possible paths.

53. (New) The method of claim 1, further comprising:

determining whether a failure has occurred that prevents the transfer of data over a first path of the plurality of possible paths;

when the failure has occurred:

failing over to one or more alternate paths when the failure has occurred;

deciding whether the first VM should be migrated to a different physical computer according to a second algorithm, which takes as inputs at least contents of the VM-specific information;

and migrating the first VM when the second algorithm returns a decision to migrate the first VM.

54. (New) The tangible computer readable medium of claim 37, wherein the method further comprises:

deciding, prior to the routing and according to a second algorithm, whether to immediately route the data transfer request or to queue the data transfer request, the second algorithm taking as inputs at least contents of the multipath routing information and contents of the VM-specific information, wherein the routing is immediately performed when the deciding results in a decision to immediately route the data transfer request.

55. (New) The tangible computer readable medium of claim 37, wherein the method further comprises:

determining whether a failure has occurred that prevents the transfer of data over a first path of the plurality of possible paths;

failing over to one or more alternate paths when the failure has occurred;

deciding, according to a second algorithm, whether the first VM should be migrated to a different physical computer system, the second algorithm taking as inputs at least whether the failure has occurred and contents of the VM-specific information; and

migrating the first VM to the different physical computer system when the second algorithm returns a decision migrate the first VM.

56. (New) The tangible computer readable medium of claim 37, wherein the VM-specific information indicates an amount of disk bandwidth that is allocated to the VM.

57. (New) The tangible computer readable medium of claim 37, wherein the multipath routing information includes a pending data transfer load for each of plurality of possible paths over which data could be routed.

58. (New) The tangible computer readable medium of claim 37, wherein:
the VM-specific information includes an identifier of the first VM as the source of the request; and

the algorithm distributes requests such that substantially all requests received from the first VM are routed over the one path, and substantially all requests from at least another VM of the plurality of VMs are routed over a second path of the plurality of possible paths.